

# On the Semantics of ReFLect as a Basis for a Reflective Theorem Prover

Tom Melham<sup>1</sup>, Raphael Cohn<sup>2</sup>, and Ian Childs<sup>1</sup>

<sup>1</sup> University of Oxford

<sup>2</sup> The Rockefeller University

**Abstract.** This paper explores the semantics of a combinatory fragment of *reFLect*, the  $\lambda$ -calculus underlying a functional language used by Intel Corporation for hardware design and verification. *ReFLect* is similar to ML, but has a primitive data type whose elements are the abstract syntax trees of *reFLect* expressions themselves. Following the LCF paradigm, this is intended to serve as the object language of a higher-order logic theorem prover for specification and reasoning—but one in which object- and meta-languages are unified. The aim is to intermix program evaluation and logical deduction through reflection mechanisms. We identify some difficulties with the semantics of *reFLect* as currently defined, and propose a minimal modification of the type system that avoids these problems.

## 1 Introduction

*ReFLect* is a strongly typed, functional programming language, similar to ML, with certain reflection features for applications in hardware design and verification [4]. The language was designed by researchers at Intel Corporation and forms the basis for Intel’s Forte formal verification environment [9]. Forte has been used at Intel to attack many challenging verification problems for real-world processor designs; one impressive achievement is the verification, without simulation, of much of the Core i7 processor execution cluster [6].

Reflection is supported through a primitive data type, *term*, whose values are *reFLect* abstract syntax trees. A *quotation*  $\langle e \rangle$  has type *term* and denotes the syntax tree of the term  $e$ . In Forte, this type is the basis for a higher-order logic theorem prover that is similar to HOL [3]. In systems like HOL, the higher-order logic ‘object language’ is built on top of the  $\lambda$ -calculus, following Church [1]. The syntax of the logic is represented by an algebraic data-type in a functional programming ‘meta-language’. *ReFLect* is designed to unify the object and meta languages in this kind of enterprise. The aim is for the *same*  $\lambda$ -calculus to be the core of both logic and meta-language. Moreover, it should be possible to move freely between *evaluation* in the interpreter and *deduction* in the theorem prover. For example, the reduction rules of *reFLect* align with inference rules of its logic, so proof can be done just by evaluation: to prove a theorem  $\langle P \rangle$ , simply strip off the quotes and check that the interpreter evaluates  $P$  to ‘true’.

The logical soundness of this depends on having the right formal semantics to justify these reflection rules. There exists an operational (i.e. reduction)

semantics for full *reFlect*, including quotation evaluation [4], but for logic a denotational semantics is needed. Krstić and Matthews have published one [7], but this omits the crucial reflection constructs that bridge logic and evaluation.

In this paper, we analyse the semantics of a simplified system, *Combinatory ReFlect*, that has some core features of full *reFlect* and includes the reflection constructs missing from the semantics in [7]. This gives a simple setting to investigate the problem while side-stepping the technicalities of variable binding in the full language. We find that even our variable-free language does not support the semantics we need, and that any reasonable logic built on *reFlect* will be inconsistent. Our proposed solution is a modification of the language’s type system that stratifies the semantics enough to avoid the problem.

## 2 Combinators with Reflection

*Combinatory ReFlect* (CR) is a variable-free system of combinatory logic that includes essentially the same reflection constructs, quotations, and typing system as full *reFlect*. Indeed, CR can be viewed as a sublanguage of full *reFlect*. The syntax of CR includes explicitly-typed versions of the combinators I, K and S, along with reflective operators *value*, *app* and *lift*. CR also supports quotation of any term  $e$  in the language, written  $\langle e \rangle$ .

The types of CR terms,  $Ty$ , are defined inductively by

$$\sigma := \text{bool} \mid \text{unit} \mid \text{term} \mid \sigma_1 \rightarrow \sigma_2$$

And the terms of CR,  $Exp$ , are given by

$$e := I_\sigma \mid K_{\sigma,\tau} \mid S_{\sigma,\tau,v} \mid \text{value}_\sigma \mid \text{lift} \mid \text{app} \mid e_1 e_2 \mid \langle e \rangle \quad (1)$$

where  $\sigma, \tau$  and  $v$  range over types. We refer to terms of the form ‘ $e_1 e_2$ ’ as *applications*, and terms of the form ‘ $\langle e \rangle$ ’ as *quotations*.

A quotation is an object-language phrase of CR that, semantically, denotes the abstract syntax tree of the term inside the quotes. A quotation  $\langle e \rangle$  is semantically different from the term  $e$ . Unquoted terms that evaluate to the same value, for example  $(S K) K e$  and  $I e$ , are semantically equal. But  $\langle (S K) K e \rangle$  and  $\langle I e \rangle$  are semantically distinct; they denote different syntax trees.

To explain the intended semantics of *value* $_\sigma$ , *lift*, and *app*, we suppose, just for a moment, that normal forms exist in some reduction system for CR. Of course we have not *established* this yet. The intended semantics, then, is that *value* $_\sigma \langle e \rangle$  should denote the normal form or evaluation of  $e$  of type  $\sigma$ . This is obviously problematic if  $e$  contains free variables, so the operational semantics of *reFlect* allows reduction of *value* $_\sigma \langle e \rangle$  only for closed  $e$  [4, p. 187]. The *lift* function reifies values into representative syntax and, in full *reFlect*, applies only to closed terms that possess a canonical ‘name’ for their value [4, §8.2]. For example, both *lift* 2 and *lift* (1 + 1) reduce to  $\langle 2 \rangle$ . In CR, we adopt an even more conservative semantics for *lift*: it applies only to terms of type *term*, and we can reduce only terms of the form *lift*  $\langle e \rangle$ . It seems obviously harmless to suppose we might have

$$\begin{array}{c}
\overline{\text{I}_\sigma : \sigma \rightarrow \sigma} \quad \overline{\text{K}_{\sigma,\tau} : \sigma \rightarrow \tau \rightarrow \sigma} \quad \overline{\text{S}_{\sigma,\tau,v} : (\sigma \rightarrow \tau \rightarrow v) \rightarrow (\sigma \rightarrow \tau) \rightarrow \sigma \rightarrow v} \\
\overline{\text{value}_\sigma : \text{term} \rightarrow \sigma} \quad \overline{\text{lift} : \text{term} \rightarrow \text{term}} \quad \overline{\text{app} : \text{term} \rightarrow \text{term} \rightarrow \text{term}} \\
\frac{e_1 : \sigma \rightarrow \tau \quad e_2 : \sigma}{e_1 e_2 : \tau} \quad \frac{e : \sigma}{\langle e \rangle : \text{term}}
\end{array}$$

**Fig. 1.** Typing rules for CR.

$$\begin{array}{c}
\text{I}_\sigma e \Rightarrow e \quad \text{K}_{\sigma,\tau} e_1 e_2 \Rightarrow e_1 \quad \text{S}_{\sigma,\tau,v} e_1 e_2 e_3 \Rightarrow e_1 e_3 (e_2 e_3) \\
\text{if } e : \sigma, \text{ then } \text{value}_\sigma \langle e \rangle \Rightarrow e \quad \text{lift} \langle s \rangle \Rightarrow \langle \langle s \rangle \rangle \\
\text{if } e_1 e_2 \text{ is well typed, then } \text{app} \langle e_1 \rangle \langle e_2 \rangle \Rightarrow \langle e_1 e_2 \rangle
\end{array}$$

**Fig. 2.** Some Reductions in CR.

a function that takes  $\langle e \rangle$  to  $\langle \langle e \rangle \rangle$ . Finally, in what follows we will need to apply one *term* to another; this is native in full *reFlect*, but in CR we have an ad-hoc combinator, **app**.

Figure 1 shows the type system of CR. The judgement  $e : \sigma$  means that the term  $e$  has type  $\sigma$ . Note that the type system does not rule out the formation of a term ‘ $\text{value}_\sigma \langle e \rangle$ ’ where  $e$  does not have type  $\sigma$ . A check for this kind of type mismatch is done only at runtime in CR, during reduction. In Figure 2 we show part of the reduction system for CR. (We omit only the rules that close  $\Rightarrow$  under reflexivity, transitivity, and congruence.)

## 2.1 Denotational Semantics

The reflective operators  $\text{value}_\sigma$  and **lift** are key components of the *reFlect* language. They serve as the essential bridge between deductive logic and program evaluation in Forte. It is these operators, omitted from the semantics in [7], that make the denotational semantics of *reFlect* difficult, as we now show.

Suppose we have a semantic function,  $\llbracket - \rrbracket \in \text{Exp} \rightarrow U$  that maps each term to an element of some universe  $U$ . A quotation  $\langle e \rangle$  represents the abstract syntax tree of its contents  $e$ . So the formal denotation of a quoted term  $\llbracket \langle e \rangle \rrbracket$ , where  $e$  is a closed term, can just be the *reFlect* term  $e$  itself, an element of  $\text{Exp} \subseteq U$ . And, indeed, this is how it is defined in [7]. With the usual semantics of function application, we would therefore expect for closed  $e : \sigma$  that

$$\llbracket \text{value}_\sigma \langle e \rangle \rrbracket = \llbracket \text{value}_\sigma \rrbracket \llbracket \langle e \rangle \rrbracket = \llbracket \text{value}_\sigma \rrbracket e = \llbracket e \rrbracket$$

So the action of the  $\text{value}_\sigma$  function seems to coincide with part of  $\llbracket - \rrbracket$  itself.

It has been shown that core  $reFIEct$ , without the reflective functions, is confluent and normalizing [7]. The authors of [7] used this normalizing result as a pre-requisite for proving the soundness of their denotational semantics. We now show that adding the reflective functions  $lift$  and  $value_\sigma$  destroys this property.

**Theorem 1.** *Combinatory  $reFIEct$  is not strongly normalizing.*

*Proof.* We proceed by constructing a term that has an infinite reduction sequence. We first define a CR expression  $f$  such that for all  $e$  of type  $term \rightarrow term$ ,  $f \langle e \rangle \Rightarrow \langle e \langle e \rangle \rangle$ . In full  $reFIEct$ ,  $f$  would be  $\lambda x. \mathbf{app} \ x \ (\mathbf{lift} \ x)$ , which in combinators is  $\mathbf{S}_{term,term,term} \ \mathbf{app} \ \mathbf{lift}$ . It is easy to check that this term is typeable with type  $term \rightarrow term$  by the rules in Figure 1. We also have the reduction sequence

$$f \langle e \rangle = \mathbf{S} \ \mathbf{app} \ \mathbf{lift} \ \langle e \rangle \Rightarrow \mathbf{app} \ \langle e \rangle \ (\mathbf{lift} \ \langle e \rangle) \Rightarrow \mathbf{app} \ \langle e \rangle \ \langle \langle e \rangle \rangle \Rightarrow \langle e \ \langle e \rangle \rangle$$

for any  $e$  of type  $term \rightarrow term$ . (We omit type subscripts for readability.)

Now, let the term  $g$  of type  $term \rightarrow term$  be  $\mathbf{S} \ (\mathbf{K} \ \mathbf{value}_{term}) \ (\mathbf{S} \ (\mathbf{K} \ f) \ \mathbf{I})$ ; for readability we omit type subscripts. In full  $reFIEct$ ,  $g$  would be  $\lambda x. \mathbf{value}_{term} \ (f \ x)$ . We immediately see that  $g \langle g \rangle$  is well typed. But now we have a circular reduction sequence:  $g \langle g \rangle \Rightarrow \mathbf{value}_{term} \ (f \ \langle g \rangle) \Rightarrow \mathbf{value}_{term} \ (\langle g \ \langle g \rangle \rangle) \Rightarrow g \langle g \rangle$ .

## 2.2 Indefinability in CR

We now show that any higher order logic built on  $reFIEct$  and containing the sublanguage CR will be inconsistent. We will proceed by supposing we do have such a logic, formulated in the usual way deriving from Church [1]. That is, we suppose a logic has been defined on top of CR in the same way that the HOL logic is defined on top of the simply-typed  $\lambda$ -calculus [3]. We then demonstrate inconsistency by a construction inspired by Tarski's Indefinability Theorem [10].

**Theorem 2.** *In a conventional higher-order logic built on CR, for any term  $\Psi : term \rightarrow bool$ , there is a typeable expression,  $\Gamma$ , such that  $\vdash \Gamma = \Psi \langle \Gamma \rangle$ .*

*Proof.* Suppose  $\Psi : term \rightarrow bool$ . Define  $\beta$  of type  $term \rightarrow bool$  such that for any  $e$  of type  $term$ ,  $\beta \ e \Rightarrow \Psi(f \ e)$ , where  $f = \mathbf{S}_{term,term,term} \ \mathbf{app} \ \mathbf{lift}$ , as in the proof of Theorem 1. In full  $reFIEct$ ,  $\beta$  would just be  $\lambda x. \Psi(f \ x)$ , where  $x$  is chosen not to occur free in  $\Psi$ . We now let  $\Gamma$  be  $\beta \ \langle \beta \rangle$  and prove  $\vdash \Gamma = \Psi \langle \Gamma \rangle$  in our assumed higher order logic:

$$\begin{aligned} \Gamma &= \beta \ \langle \beta \rangle && \text{--definition of } \Gamma \\ &= \Psi(f \ \langle \beta \rangle) && \text{--reduction of } \beta \ \langle \beta \rangle \\ &= \Psi \ \langle \beta \ \langle \beta \rangle \rangle && \text{--reduction of } f \ \langle \beta \rangle \\ &= \Psi \ \langle \Gamma \rangle && \text{--definition of } \Gamma \end{aligned}$$

**Corollary 1.** *Any conventional higher order logic built on CR is inconsistent.*

*Proof.* The proof is straightforward and we sketch it here. As discussed, CR contains a truth predicate,  $value_{bool}$ . We then have a falsity predicate  $isfalse$  defined to be  $\mathbf{S} \ (\mathbf{K} \ \neg) \ (\mathbf{S} \ (\mathbf{K} \ \mathbf{value}_{bool}) \ \mathbf{I})$ . Taking  $\Psi$  in Theorem 2 to be  $isfalse$ , we conclude that any logic built on CR can prove  $\vdash \Gamma = isfalse \ \langle \Gamma \rangle$ . But then  $value_{bool} \ \langle \Gamma \rangle = \Gamma = isfalse \ \langle \Gamma \rangle = \neg (value_{bool} \ \langle \Gamma \rangle)$ , a contradiction.

### 3 Stratified Typing for Terms

Analysis of the above results for CR reveals that its fundamental flaw is that all quotations have a single type, *term*. This allows the circularities that prevent normalization and make it logically inconsistent. We now sketch a variant, *Stratified Combinatory ReFlect* (SCR), that avoids these pitfalls. The basic idea is simple: the type of a quotation will carry with it the type of the term inside.

The syntax and reduction semantics of SCR are essentially the same as in CR. The difference is in the types. The types of SCR terms are defined by

$$\sigma := \text{unit} \mid (\sigma)\text{term} \mid \sigma_1 \rightarrow \sigma_2$$

and the typing rules for the reflective combinators and quotations are

$$\frac{}{\text{value}_\sigma : (\sigma)\text{term} \rightarrow \sigma} \quad \frac{}{\text{lift}_\sigma : (\sigma)\text{term} \rightarrow ((\sigma)\text{term})\text{term}}$$

$$\frac{}{\text{app} : (\sigma \rightarrow \tau)\text{term} \rightarrow (\sigma)\text{term} \rightarrow (\tau)\text{term}} \quad \frac{e : \sigma}{\langle e \rangle : (\sigma)\text{term}}$$

The other combinators and applications are typed as they are in CR.

This new typing schemes rules out applying a function to its quoted self. It bans the self-application used to define the nonterminating ‘ $g \langle g \rangle$ ’ in Theorem 1, and it rules out the ‘ $\beta \langle \beta \rangle$ ’ construction used in the proof of Theorem 2.

### 4 Conclusions and Discussion

This paper has explored the semantics of a reflective combinatory logic that shares key features with *reFlect*, a functional language intended to unify computation and deduction in a practical engineering setting. Our analysis suggests that the type of quotations in *reFlect* must be parameterized by the type of their contents. We speculate that, with this adjustment to the type system, a set-theoretic semantics of full *reFlect*, including reflection, will be possible.

To make *ReFlect* attractive to verification engineers, it has a simple Hindley-Milner type system. This means quotations cannot have types of the form  $(\sigma)\text{term}$  without making the definitions of certain common functions over terms untypeable. For example, we cannot define `operand`  $\langle e_1 e_2 \rangle = \langle e_1 \rangle$ , since this function would have to have an existential type  $(\alpha)\text{term} \rightarrow \exists \beta. (\beta \rightarrow \alpha)\text{term}$ . Also problematic are functions defined recursively over the syntax of terms, which are ubiquitous in theorem-prover code.

Some developments in functional programming, subsequent to the design of *reFlect*, offer a way forward. *Generalized algebraic data types* (GADTs) are a generalization of standard algebraic data types that take a modest step towards dependent types [8]. GADTs allow for algebraic data type constructors to have parameters that can be instantiated to specific types within the body of a function defined over values of the type. We speculate that one could treat quotations and the type  $(\sigma)\text{term}$  as a GADT in *reFlect*, with the aim of making it possible to define the term-traversing functions needed to implement a theorem prover.

*reF $\mathcal{L}$ ect* currently includes neither GADTs nor  $(\sigma)$ term. A major project for the future is to develop a full reflective language with GADTs and a theorem prover based on the semantic insights in this paper. A clear first step is to investigate whether GADTs, as found for example in Haskell, are indeed sufficient to develop a theorem prover with a parameterized type  $(\sigma)$ term of terms. We have done an initial investigation of this, and the answer seems to be ‘partly’; the main difficulty seems to be finding a satisfactory treatment of polymorphism.

We are of course aware that there are functional languages and logical calculi with more flexible type systems, among them System F, Coq, and HOL-Omega. But in this industrially-motivated work we have been exploring options that give practicing engineers as simple and intuitive a functional programming language as possible, and so aim to remain close to the Hindley-Milner type system.

More generally, fast object-language evaluation has been a goal of theorem prover designers since the earliest days. Approaches include term data structures that optimise symbolic evaluation by proof, and extraction of programs in standard functional languages [2]. The work presented here is distinguished in attempting direct unification of object- and meta-languages, of symbolic reasoning and direct program execution.

*Acknowledgment* This work leans heavily on a preliminary sketch of the semantics of full *reF $\mathcal{L}$ ect* devised in collaboration with Jim Grundy and Sava Krstić [5]

## References

1. A. Church. A formulation of the simple theory of types. *The Journal of Symbolic Logic*, 5:56–68, 1940.
2. Coq Development Team. *The Coq Proof Assistant: Reference Manual, V8.4*. Inria, Aug. 2012.
3. M. J. C. Gordon and T. F. Melham, editors. *Introduction to HOL: A Theorem Proving Environment for Higher Order Logic*. CUP, 1993.
4. J. Grundy, T. Melham, and J. O’Leary. A reflective functional language for hardware design and theorem proving. *JFP*, 16(2):157–196, Mar. 2006.
5. J. Grundy, T. F. Melham, and S. Krstic. Towards a denotational semantics of *reF $\mathcal{L}$ ect*. Unpublished presentation slides, 2007.
6. R. Kaivola, R. Ghughal, N. Narasimhan, A. Telfer, J. Whittemore, S. Pandav, A. Slobodová, C. Taylor, V. Frolov, E. Reeber, and A. Naik. Replacing testing with formal verification in Intel Core i7 processor execution engine validation. In *Computer Aided Verification, CAV 2009*, pages 414–429. Springer, 2009.
7. S. Krstić and J. Matthews. Semantics of the *reF $\mathcal{L}$ ect* language. In E. Moggi and D. S. Warren, editors, *PPDP*, pages 32–42. ACM, 2004.
8. S. L. Peyton Jones, D. Vytiniotis, S. Weirich, and G. Washburn. Simple unification-based type inference for GADTs. In *ICFP*, pages 50–61. ACM, 2006.
9. C.-J. H. Seger, R. B. Jones, J. W. O’Leary, T. Melham, M. D. Aagaard, C. Barrett, and D. Syme. An industrially effective environment for formal hardware verification. *IEEE TCAD*, 24(9):1381–1405, Sept. 2005.
10. A. Tarski. The concept of truth in formalized languages. In J. Corcoran, editor, *Logic, Semantics, Metamathematics*, pages 152–278. Hackett Publishing Co., 1983.